



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

APPLICATION NO.	FILING DATE	FIRST NAMED INVENTOR	ATTORNEY DOCKET NO.	CONFIRMATION NO.
10/077,345	02/15/2002	Roger L. Haskin	POU920020010US1	2492
23334	7590	01/06/2005	EXAMINER	
FLEIT, KAIN, GIBBONS, GUTMAN, BONGINI & BIANCO P.L. ONE BOCA COMMERCE CENTER 551 NORTHWEST 77TH STREET, SUITE 111 BOCA RATON, FL 33487			CHANNAVAJJALA, SRIRAMA T	
		ART UNIT		PAPER NUMBER
		2164		
DATE MAILED: 01/06/2005				

Please find below and/or attached an Office communication concerning this application or proceeding.

Office Action Summary	Application No.	Applicant(s)	
	10/077,345	HASKIN ET AL.	
	Examiner	Art Unit	
	Srirama Channavajjala	2164	

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

Period for Reply

A SHORTENED STATUTORY PERIOD FOR REPLY IS SET TO EXPIRE 3 MONTH(S) FROM THE MAILING DATE OF THIS COMMUNICATION.

- Extensions of time may be available under the provisions of 37 CFR 1.136(a). In no event, however, may a reply be timely filed after SIX (6) MONTHS from the mailing date of this communication.
- If the period for reply specified above is less than thirty (30) days, a reply within the statutory minimum of thirty (30) days will be considered timely.
- If NO period for reply is specified above, the maximum statutory period will apply and will expire SIX (6) MONTHS from the mailing date of this communication.
- Failure to reply within the set or extended period for reply will, by statute, cause the application to become ABANDONED (35 U.S.C. § 133). Any reply received by the Office later than three months after the mailing date of this communication, even if timely filed, may reduce any earned patent term adjustment. See 37 CFR 1.704(b).

Status

1) Responsive to communication(s) filed on 26 January 2004.

2a) This action is **FINAL**. 2b) This action is non-final.

3) Since this application is in condition for allowance except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte Quayle*, 1935 C.D. 11, 453 O.G. 213.

Disposition of Claims

4) Claim(s) 1-26 is/are pending in the application.

4a) Of the above claim(s) _____ is/are withdrawn from consideration.

5) Claim(s) _____ is/are allowed.

6) Claim(s) 1-26 is/are rejected.

7) Claim(s) _____ is/are objected to.

8) Claim(s) _____ are subject to restriction and/or election requirement.

Application Papers

9) The specification is objected to by the Examiner.

10) The drawing(s) filed on _____ is/are: a) accepted or b) objected to by the Examiner.
Applicant may not request that any objection to the drawing(s) be held in abeyance. See 37 CFR 1.85(a).
Replacement drawing sheet(s) including the correction is required if the drawing(s) is objected to. See 37 CFR 1.121(d).

11) The oath or declaration is objected to by the Examiner. Note the attached Office Action or form PTO-152.

Priority under 35 U.S.C. § 119

12) Acknowledgment is made of a claim for foreign priority under 35 U.S.C. § 119(a)-(d) or (f).

a) All b) Some * c) None of:
 1. Certified copies of the priority documents have been received.
 2. Certified copies of the priority documents have been received in Application No. _____.
 3. Copies of the certified copies of the priority documents have been received in this National Stage application from the International Bureau (PCT Rule 17.2(a)).

* See the attached detailed Office action for a list of the certified copies not received.

Attachment(s)

1) Notice of References Cited (PTO-892)
 2) Notice of Draftsperson's Patent Drawing Review (PTO-948)
 3) Information Disclosure Statement(s) (PTO-1449 or PTO/SB/08)
 Paper No(s)/Mail Date _____

4) Interview Summary (PTO-413)
 Paper No(s)/Mail Date. _____

5) Notice of Informal Patent Application (PTO-152)
 6) Other: _____

DETAILED ACTION

1. Examiner acknowledges applicant's amendment filed on 10/1/2004.
2. Claims 1-26 have been amended [10/1/2004].
3. Claims 1-26 are pending in this application.
4. In view of the applicant submitted "terminal disclaimer", rejection under obviousness-type double patenting as set forth in the previous office action is hereby withdrawn.

Drawings

5. The Drawings filed on 6/3/2002 are acceptable for examination purpose.

Information Disclosure Statement

6. The information disclosure statement filed on 1/26/2004, paper no. # 4 is in compliance with the provisions of 37 CFR 1.97, and has been considered and a copy was enclosed with this Office Action. [see paper no. # 5].

Claim Rejections - 35 USC § 102

The following is a quotation of the appropriate paragraphs of 35 U.S.C. 102 that form the basis for the rejections under this section made in this Office action:

A person shall be entitled to a patent unless –

(e) the invention was described in (1) an application for patent, published under section 122(b), by another filed in the United States before the invention by the applicant for patent or (2) a patent granted on an application for patent by another filed in the United States before the invention by the applicant for patent, except that an international application filed under the treaty defined in section 351(a) shall have the effects for purposes of this subsection of an application filed in the United States only if the international application designated the United States and was published under Article 21(2) of such treaty in the English language.

7. Claim 1-25 are rejected under 35 U.S.C. 102(e) as being anticipated by Kazar et al., [hereafter Kazar], US Pub.No. 2002/0112022.
8. As to Claim 1,7, 13, 19,Kazar teaches a system which including 'providing a file system snapshot' [see page 6, col 2, 0100], file system snapshot corresponds to Kazar's file system snapshot as described in page 6, 0100, further it is also noted that Kazar specifically directed to volume replication for making clone that creates snapshot of the volume as detailed in page 5, col 1, 0082; 'generating a snapshot dataset for a source file in a file system, wherein the snapshot dataset contains substantially no data and no metadata [page 5, col 1, 0082,0084, fig 1], source file in a file system corresponds to fig 1, NFS server, generating snapshot dataset corresponds to creating a point in time snapshot of that volume as detailed in page 5, col 1, 0082, snapshot dataset contains substantially no data and no metadata corresponds to propagating changes or update changes of a

clone to remote sites by only sending the data blocks that have changed since the last replica was propagated as detailed in page 5, col 1, 0084;

'copying into a first inode within the snapshot dataset in response to only modifying metadata of the source file, at least a portion of metadata within a second inode corresponding to the source file, [page 1, col 2, 0020-0021, page 4, col 1, 0073, page 5, col 1, 0085, fig 1, fig 8-9], Kazar is directed to vnode operations layer, more specifically vnode, inode operations in a file system structure as detailed in fig 1, Kazar also specifically directed to copy-on-write operations to create cloned inodes [see page 1, col 2, 0020],; first inode, second inode, corresponds to fig 2, inode 1, inode 2;

'storing, into the first inode,[see page 1, col 2, 0017-0018], Kazar specifically teaches storage layer underlying the block and file level servers that performs data management operations as detailed in fig 11;

'disk address values equal to a ditto address to indicate that the disk address is an invalid disk address' [page 3, col 2, 0068, 0069, page 4, col 1, 0070], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status information about a file [see page 3, col 2, 0068, line 1-3], Kazar also teaches clone volume is a copy , it has all the information or retains all the properties of original snapshot volume, further it shares all the disk space with the original volume, when clone is created, each file inode is copied, with the result that the copied inode points to the same data blocks as the original that corresponds to disk address values and content equal to a ditto address to indicate that the disk address,

furthermore, Kazar also suggested "if the addressed disk block is an indirect block, all deeper disk blocks are also shared as well" [see page 4, col 1, 0071], therefore, if the disk address is an invalid disk address, it does share other disk blocks, and the status is ensured whether invalid disk address, whether disk block is freed and like

disk address of a data block corresponding to disk block addresses that are associated with bit information of copy on write operations [see page 4, col 1, 0071]

As best understood by the examiner, shadow inode is created when a file is opened, in other words, when copy-on write operations are performed, further, a copy is made of the original inode, pointing to the same data and indirect blocks, therefore, the shadow thus refers to the same data as the original, but is not yet referred to by any directory. It is however, noted that operations on the file use the shadow inode because each time the file is modified, or updated, a copy is made of the target block, further, the copy replaces the original in the shadow version of the file and modifications are made to the copy of the file. In order to make sure each respective block is copied, a copy on write bit is associated with each block pointer whether direct or indirect [see Kazar: page 1, col 2, 0021]. Because copy on write bit is associated with each data block pointer, it is normally set to zero, however, it may bit may also set to "1" when the data block is modified or copied, further if a change is made to specific data block and respective copy on write bit is already set to "1", no further copying is necessary [see page 4, col 1, 0073]. It is however, should be noted that this applies to indirect blocks, when they are modified, a copy is made for the shadow version of the file

9. As to Claim 2, 8, 14, 20, most of the limitations of this claim have been noted in the rejection of Claim 1 above. In addition, with respect to the claimed feature Kazar disclosed 'copying to the first inode in the snapshot dataset, in response to only appending to the source file, at least a portion of metadata within the second inode corresponding to the source file, [page 1, col 2, 0020-0021, page 4, col 1, 0073, page 5, col 1, 0085, fig 1, fig 8-9], Kazar suggests data blocks contain a copy-tree-on write related to another clone inode, that corresponds to copying to the first inode in the snapshot dataset;

'storing, into the first inode, [see page 1, col 2, 0017-0018], Kazar specifically teaches storage layer underlying the block and file level servers that performs data management operations as detailed in fig 11;

'disk address values equal to a ditto address to indicate that the disk address is an invalid disk address' [page 3, col 2, 0068, 0069, page 4, col 1, 0070], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status information about a file [see page 3, col 2, 0068, line 1-3], Kazar also teaches clone volume is a copy, it has all the information or retains all the properties of original snapshot volume, further it shares all the disk space with the original volume, when clone is created, each file inode is copied, with the result that the copied inode points to the same data blocks as the original that corresponds to disk address values and content equal to a ditto address to indicate that the disk address, furthermore, Kazar also suggested "if the addressed disk block is an indirect block, all deeper disk blocks are also shared as well" [see page 4, col 1, 0071], therefore, if the

disk address is an invalid disk address, it does share other disk blocks, and the status is ensured whether invalid disk address, whether disk block is freed and like

disk address of a data block corresponding to disk block addresses that are associated with bit information of copy on write operations [see page 4, col 1, 0071]

10. As to Claim 3, 9, 15, most of the limitations of this claim have been noted in the rejection of Claim 2 above. In addition, with respect to the claimed feature Kazar disclosed 'copying to the first inode in the snapshot dataset second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, when the data block corresponding to the source file is overwritten or deleted, wherein the first inode includes disk address of the data block which was written in the snapshot dataset' [page 4, col 1, 0074, col 2, 0079-0080], Kazar specifically teaches file delete operations does all disk blocks in the indirect block tree associated with the file being deleted as detailed in page 4, col 2, 0079, first inode, second inode, corresponds to fig 2, inode 1, inode 2;

11. As to Claim 4, 10, 16, most of the limitations of this claim have been noted in the rejection of Claim 3 above. In addition, with respect to the claimed feature Kazar disclosed 'accessing the first inode of the snapshot dataset corresponding to the source file'[page 2, col 1, 0022], first inode corresponds to fig 2, inode 1; 'determining whether the first inode includes a valid disk address' [page 3, col 2, 0060,0068];

'wherein if the first inode includes a valid disk address, then reading a data block referenced by the disk address' [page 3, col 2, 0068];

'wherein if the shadow inode contains the ditto address, then retrieving the second inode of the source file and retrieving a data block referenced by a disk address in the second inode of the source file' [page 3, col 2, 0068,0069, page 4, col 1, 0071]

12. As to Claim 5, 11,17, most of the limitations of this claim have been noted in the rejection of Claim 3 above. In addition, with respect to the claimed feature Kazar disclosed 'copying to the first inode in the snapshot dataset the metadata within the second inode corresponding to the source file and copying to the snapshot dataset the data block corresponding to the source file, when the data block corresponding to the source file is overwritten or deleted, wherein the first inode includes disk address of the data block which was written in the snapshot dataset' [page 4, col 1, 0074, col 2, 0079-0080], Kazar specifically teaches file delete operations does all disk blocks in the indirect block tree associated with the file being deleted as detailed in page 4, col 2, 0079; 'wherein the indirect block includes a disk address of at least one data block which was written in the snapshot dataset' [page 1, col 2, 0021].

13. As to Claim 6, 12, 18,most of the limitations of this claim have been noted in the rejection of Claim 5 above. In addition, with respect to the claimed feature Kazar disclosed 'accessing the first inode corresponding to the source file'[page 2, col 1, 0022];

'determining whether the first inode includes a disk address' [page 3, col 2, 0060,0068]; 'wherein if the shadow inode includes a valid disk address, then retrieving an indirect block referenced by the valid disk address and at least one data block defined by at least one disk address in the indirect block' [page 1, col 2,0021, page 3, col 2, 0068];

'wherein if the first inode does not include a valid disk address, retrieving the second inode of the source file, then retrieving an indirect block referenced by a disk address in the second inode of the source file and retrieving at least one data block referenced by at least one disk address in the indirect block' [page 4, col 1, 0071, page 5, col 1, 0083].

14. As to Claim 21,most of the limitations of this claim have been noted in the rejection of Claim 20 above. In addition, with respect to the claimed feature Kazar disclosed 'a data block corresponding to the source file in the snapshot dataset, wherein the data block is copied to the snapshot dataset when the original data block is overwritten' [page 4, col 1, 0070, 0075];

'a first inode in the snapshot dataset, the first inode containing metadata copied from an inode in the source file, wherein the first inode is generated when the data block corresponding to the source file is overwritten or deleted and wherein the first inode includes a disk address of the data block which was written in the snapshot dataset'[page 4, col 1, 0074, col 2, 0079-0080].

15. As to Claim 22, most of the limitations of this claim have been noted in the rejection of Claim 21 above. In addition, with respect to the claimed feature Kazar disclosed 'a first inode in a snapshot dataset, the first inode corresponding to a source file' [page 2, col 1, 0025, fig 2];

'a ditto address value stored in the first inode to indicate an invalid disk address, [page 3, col 2, 0068,0069, page 4, col 1, 0070], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status information about a file [see page 3, col 2, 0068, line 1-3], Kazar also teaches clone volume is a copy , it has all the information or retains all the properties of original snapshot volume, further it shares all the disk space with the original volume, when clone is created, each file inode is copied, with the result that the copied inode points to the same data blocks as the original that corresponds to disk address values and content equal to a ditto address to indicate that the disk address, furthermore, Kazar also suggested "if the addressed disk block is an indirect block, all deeper disk blocks are also shared as well' [see page 4, col 1, 0071], therefore, if the disk address is an invalid disk address, it does share other disk blocks, and the status is ensured whether invalid disk address, whether disk block is freed and like [page 3, col 2, 0068], Kazar specifically suggests inode points to a data blocks by gibing their address as detailed in 0068, line 1-2;

'an inode of the source file referencing the daa block' [see fig 1-2];

16. As to Claim 23, most of the limitations of this claim have been noted in the rejection of Claim 21 above. In addition, with respect to the claimed feature Kazar disclosed a first inode in a snapshot dataset, the first inode corresponding to an indirect block within 'source file' [page 2, col 1, 0025];

'a disk address included in the shadow inode' [page 3, col 2, 0068], Kazar specifically suggests inode points to a data blocks by gibing their address as detailed in 0068, line 1-2;

'a ditto address value stored in the first inode to indicate an invalid disk address, [page 3, col 2, 0068,0069, page 4, col 1, 0070], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status information about a file [see page 3, col 2, 0068, line 1-3], Kazar also teaches clone volume is a copy , it has all the information or retains all the properties of original snapshot volume, further it shares all the disk space with the original volume, when clone is created, each file inode is copied, with the result that the copied inode points to the same data blocks as the original that corresponds to disk address values and content equal to a ditto address to indicate that the disk address, furthermore, Kazar also suggested "if the addressed disk block is an indirect block, all deeper disk blocks are also shared as well" [see page 4, col 1, 0071], therefore, if the disk address is an invalid disk address, it does share other disk blocks, and the status is ensured whether invalid disk address, whether disk block is freed and like [page 3, col 2, 0068], Kazar specifically suggests inode points to a data blocks by gibing their address as detailed in 0068, line 1-2;

'an inode of the source file referencing the indirect block' [see fig 1-2];

17. As to Claim 24-25, Kazar teaches a system which including 'determining the existence of an older snapshot' [page 5, col 1, 0081];

'wherein if there is an older snapshot, determining the existence of a ditto address in the older snapshot in an inode or a data block in the first snapshot wherein the ditto address indicates an invalid disk address [page 3, col 2, 0068,0069, page 4, col 1, 0070], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status information about a file [see page 3, col 2, 0068, line 1-3], Kazar also teaches clone volume is a copy , it has all the information or retains all the properties of original snapshot volume, further it shares all the disk space with the original volume, when clone is created, each file inode is copied, with the result that the copied inode points to the same data blocks as the original that corresponds to disk address values and content equal to a ditto address to indicate that the disk address, furthermore, Kazar also suggested "if the addressed disk block is an indirect block, all deeper disk blocks are also shared as well' [see page 4, col 1, 0071], therefore, if the disk address is an invalid disk address, it does share other disk blocks, and the status is ensured whether invalid disk address, whether disk block is freed and like [page 3, col 2, 0068], Kazar specifically suggests inode points to a data blocks by gibing their address as detailed in 0068, line 1-2;

'wherein if there is no older snapshot, deleting any inode or data block in the first snapshot' [page 4, col 2, 0079].

18. Claim 26 is rejected under 35 U.S.C. 102(e) as being anticipated by Lewis et al., [hereafter Lewis], US Pub.No. 2002/0083037.

19. As to Claim 26, Lewis teaches a system which including 'wherein if there is a most recent snapshot, the most recent snapshot not being the first snapshot, copying to the most recent snapshot any inode data block in the file system referenced by the most recent snapshot which shall be modified by the restoration of the first snapshot' [see Abstract, col 2, especially line 26-43], most recent snapshot corresponds to Lewis's most recently created snapshot;

'wherein if there is an inode or a data block in the first snapshot, copying the inode or data block in the first snapshot to the file system' [page 4, col 1, 0058, page 3, col 2, 0050, page 5, col 2, 0096];

'wherein if there is a ditto disk address in the first snapshot, wherein the ditto address indicates an invalid disk address, copying to the filesystem the inode or data block of the most recent snapshot that corresponds to an inode with the ditto disk address and that contains a valid disk address [page 4, col 2, 0063-0064], Lewis specifically teaches each snapshot includes all the information related to block and is equivalent to older block from a previous active file system that corresponds to copying exactly or ditto information especially related to snapshot.

Response to Arguments

20. Applicant's arguments filed on 10/01/2004 with respect to claims 1-26 have been fully considered but they are not persuasive, for examineer's response, see discussion below.

a) At page 15, claims 1,7,13,19, applicant argues that the present invention overcomes this requirement to create a new inode clone by "generating a snapshot dataset for a source file in a file system, wherein the snapshot dataset contains substantially no data and no metadata" as claimed

As to the above argument [a] as best understood by the examiner, Kazar disclosed 'generating a snapshot dataset for a source file in a file system, wherein the snapshot dataset contains substantially no data and no metadata [page 5, col 1, 0082,0084, fig 1], source file in a file system corresponds to fig 1, NFS server, generating snapshot dataset corresponds to creating a point in time snapshot of that volume as detailed in page 5, col 1, 0082, snapshot dataset contains substantially no data and no metadata corresponds to propagating changes or update changes of a clone to remote sites by only sending the data blocks that have changed since the last replica was propagated as detailed in page 5, col 1, 0084;

Art Unit: 2164

b) At page 16, claims 1,7,13,19, applicant assert that the elements [copying into a first inode within the snapshot dataset,.....] of these limitation, which comprise not including any disk address values of data blocks corresponding to the source file, and storing disk address values equal to a “ditto” address, are not taught or suggested by the prior art of record.

c) At page 16, claims 1,7,13,19, applicant argues that amended independent claims, at least a portion of the metadata within the inode corresponding to the source data, which is metadata in the inode referred to as the second inode, is copied into the inode of the snapshot'

As to the above argument [b-c], Kazar teaches these limitations particularly 'copying into a first inode within the snapshot dataset in response to only modifying metadata of the source file, at least a portion of metadata within a second inode corresponding to the source file, [page 1, col 2, 0020-0021, page 4, col 1, 0073, page 5, col 1, 0085, fig 1,fig 8-9], Kazar is directed to vnode operations layer, more specifically vnode, inode operations in a file system structure as detailed in fig 1, Kazar also specifically directed to copy-on-write operations to create cloned inodes [see page 1, col 2, 0020],; first inode, second inode, corresponds to fig 2, inode 1, inode 2; 'storing, into the first inode,[see page 1, col 2, 0017-0018], Kazar specifically teaches storage layer underlying the block and file level servers that performs data management operations as detailed in fig 11;

'disk address values equal to a ditto address to indicate that the disk address is an invalid disk address' [page 3, col 2, 0068,0069, page 4, col 1, 0070], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status information about a file [see page 3, col 2, 0068, line 1-3], Kazar also teaches clone volume is a copy , it has all the information or retains all the properties of original snapshot volume, further it shares all the disk space with the original volume, when clone is created, each file inode is copied, with the result that the copied inode points to the same data blocks as the original that corresponds to disk address values and content equal to a ditto address to indicate that the disk address, furthermore, Kazar also suggested "if the addressed disk block is an indirect block, all deeper disk blocks are also shared as well" [see page 4, col 1, 0071], therefore, if the disk address is an invalid disk address, it does share with other disk blocks, and the status is ensured whether invalid disk address, whether disk block is freed and like

d) At page 16, claims 1,7,13,19, applicant argues that the present invention stores "ditto" values as disk address values in the inodes of the snapshot dataset and not the physical address of the physical data blocks themselves.

As to the argument [d], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status information about a file [see page 3, col 2, 0068, line 1-3,page 4, col 1, 0071], also Kazar teaches copy tree on write or CTW setting specific value for the block address by the pointer is

effective to snapshot or clones, hence exact or ditto values as disk address values in data block and associated snapshot are maintained.

e) At page 17, claims 2,8,14,20, applicant argues that Kazar simply teaches copying inodes and does not teach or discuss selectively copying data from a source file system inode to a clone or replica file system inode in response to a process, as is recited for these amended dependent claims'

As to the argument [e], as best understood by the examiner, Kazar teaches copy-tree-on-write operation that specifically selects block pointed inode to clone inode block that corresponds to copying information or data from source inode to clone inode as detailed in page 1, col 2, 0021.

f) At page 17, claims 2,8,14,20, applicant argues that the use of a "ditto" address or its equivalent in the newly created inode, which is defined in these amended claims as indicating that the disk address is an invalid disk address is further not taught or suggested in the prior art of record.

As to the argument [f], as best understood by the examiner, disk address values equal to a ditto address to indicate that the disk address is an invalid disk address' [page 3, col 2, 0068,0069, page 4, col 1, 0070], Kazar specifically teaches inode points, data blocks and respective data blocks address, further inode also have specific status

information about a file [see page 3, col 2, 0068, line 1-3], Kazar also teaches clone volume is a copy , it has all the information or retains all the properties of original snapshot volume, further it shares all the disk space with the original volume, when clone is created, each file inode is copied, with the result that the copied inode points to the same data blocks as the original that corresponds to disk address values and content equal to a ditto address to indicate that the disk address, furthermore, Kazar also suggested "if the addressed disk block is an indirect block, all deeper disk blocks are also shared as well" [see page 4, col 1, 0071], therefore, if the disk address is an invalid disk address, it does share other disk blocks, and the status is ensured whether invalid disk address, whether disk block is freed and like.

g) At page 17-18, claims 4,10,16, applicant argues that the since only one inode points to a particular data block, and "ditto" addresses are used to indicate that the real disk address of a data block must be found elsewhere, as is recited for amended claims 4,10,16.

As to the above argument [g], Kazar specifically teaches when snapshot is created, each inode and associated files are copied, these copied inode points to the same data blocks as the original that corresponds to exact copy or "ditto" of a particular data block and ditto addresses are used and same information are part of the snapshot,[page 4, col 1, 0069-0070].

h) At page 18, claims 6,12,18, these claims recite similar features and benefits, and similarly distinguish over the cited prior art as claims 4,10,16.

As to the argument [h], examiner applies above discussed arguments to 6,12,18.

i) At page 18, claims 22-23, applicant argues that the ditto address is specified in amended claims 22 and 23 to indicate an invalid disk address. As discussed above, Kazar and the other prior art of record is silent as to an inode.....

As to the above argument [i], examiner rejected amended claims 22-23 as detailed above.

j) At page 18-19, claims 24-25, applicant argues that claims 24 and 25 to more clearly specify the characteristics of the snapshots.....including ditto addresses which are discussed

As to the above argument, examiner rejected amended claims 24-25 as detailed above.

k) At page 19, claim 26, applicant argues that Although the system of Lewis uses a copy-on-write mechanism, Lewis is silent as to any form of “logical addressing”, such as through the use of “ditto” addresses as is claimed for the present invention.

As to the above argument [k], Lewis specifically teaches each snapshot includes all the information related to block and is equivalent to older block from a previous active file system that corresponds to copying exactly or ditto information especially related to snapshot [page 4, col 2, 0063-0064].

Conclusion

The prior art made of record

- a. US Pub. No. 2002/0112022
- b. US Pub.No. 2002/0083037

The prior art made of record and not relied upon is considered pertinent to applicant's disclosure

- c. US Patent No. 6341341
- d. US Patent No. 6205450
- e. US Patent No. 5764972
- f. US Patent No. 6038639
- g. US Patent No. 6654912
- h. US Patent No. 6173293
- j. US Pub. 2003/0158873
- k. US Pub. 2003/0158862
- l. US Pub. 2003/0140204

- m. US Patent No. 6484186
- n. US Patent No. 5991771
- o. US Patent No. 5678042
- p. US Pub. 2003/0140070
- q. Douglas et al., « Deciding when to forget in the elephant file syste, 17th ACM Symposium on operating systems principles SOSP, 1999 ACM pp110-123
- r. Vitor Santos Costa, LIACC & DCC-FCUP, "COWL: Copy-on-write for logic programs pp 1-8
- s. HITACHI data systems, "Hitachi quickshadow copy-on-write snapshot software, 2004 4 pages
- t. LSI LOGIC STORAGE SYSTEMS "snapshot feature" © 2002 2 pages
- u. Snap Appliance Technology Brief, © 2004 rev.2 4 pages
- v. Gregory R G et al. Embedded Inodes and explicit grouping: exploiting disk bandwith for small files, proceedings of the USENIX 1997 annual technical conference, Jan 1997, 17 pages
- w. Gregory Ganger, Soft updates: a solution to the metadata update problem in file system, pp 1-44.
- x. Darren EV et al. A holesome file system, submitted to USENIX 1996, pp 1-10

THIS ACTION IS MADE FINAL. Applicant is reminded of the extension of time policy as set forth in 37 CFR 1.136(a).

A shortened statutory period for reply to this final action is set to expire THREE MONTHS from the mailing date of this action. In the event a first reply is filed within TWO MONTHS of the mailing date of this final action and the advisory action is not mailed until after the end of the THREE-MONTH shortened statutory period, then the shortened statutory period will expire on the date the advisory action is mailed, and any extension fee pursuant to 37 CFR 1.136(a) will be calculated from the mailing date of the advisory action. In no event, however, will the statutory period for reply expire later than SIX MONTHS from the mailing date of this final action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to Srirama Channavajjala whose telephone number is 571-272-4108. The examiner can normally be reached on Monday-Friday from 8:00 AM to 5:30 PM Eastern Time.

If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Dov Popvici, can be reached on 571-272-4083. The fax phone numbers for the organization where the application or proceeding is assigned is 703/872-9306

Information regarding the status of an application may be obtained from the Patent Application Information Retrieval (PAIR) system. Status information for published applications may be obtained from either Private PAIR or Public PAIR. Status information for unpublished applications is available through Private PAIR only. For more information about the PAIR system, see <http://pair-direct.uspto.gov>. Should you have questions on access to the Private PAIR system, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free)

sc
Patent Examiner
December 29, 2004.


SRIRAMA CHANNAVAJJALA
PRIMARY EXAMINER